UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/825,539	04/15/2004	James D. Ballew	064747.1014	8244

45507 7590 08/09/2010

BAKER BOTTS LLP
2001 ROSS AVENUE
6TH FLOOR
DALLAS, TX 75201-2980

EXAMINER	
GREENE, JOSEPH L	

ART UNIT	PAPER NUMBER
2451	

NOTIFICATION DATE	DELIVERY MODE
08/09/2010	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

PTOmail3@bakerbotts.com
PTOmail4@bakerbotts.com

Office Action Summary	Application No.	Applicant(s)	
	10/825,539	BALLEW ET AL.	
	Examiner	Art Unit	
	JOSEPH GREENE	2451	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *27 May 2010*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-39* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-39* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *03/26/2008, 05/27/2010, 07/29/2010*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

Conclusion

1. Claims 1 – 39 are currently pending in this application.

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

> Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

2. Claims 1-11 and 37 are rejected under 35 U.S.C. 101 as the claimed invention is directed to non-statutory subject matter.

3. With respect to claims 1 and 37, they are directed towards "**software residing on one or more computer systems.**" While the claims do mention items that are to be interpreted as hardware, within the nodes, the system is directed specifically towards the software modules that are claimed in the preamble. The mentioning of nodes that contain hardware, in the system, are solely directed towards the intended use of the software modules. Thus, the system is directed towards software per se. Furthermore, claims 2-12 are dependent upon claim 1 and are directed towards the same software modules. Therefore, they are also rejected.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

5. Claims 1-3, 9, 12-14, 23-25, and 34-36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blumrich et al. (Pre-Grant Publication No. US 2004/0103218 A1), hereinafter Blumrich, Prael et al. (Patent No. US 7,065,764 B1), hereinafter Prael.

6. With respect to claims 1 and 23, Blumrich disclosed software residing at one or more computer systems collectively operable to execute the software, the software comprising (0024, lines 28-33 and 0025, lines 4-7): a plurality of cluster agents, each cluster agent associated with one of a plurality of nodes, each node comprising a switching fabric and at least two processors integrated to the card ([0057], lines 1-14), the cluster agent operable to determine a status of the associated node ([0210], lines 4-9); However, Blumrich did not explicitly state a cluster management engine communicably coupled to the plurality of nodes and operable to dynamically allocate a particular subset of the plurality of nodes to a particular job based on the determined status of each of one or more of the plurality of nodes and execute the job selected from a queue comprising a plurality of jobs using the particular subset. Also, Blumrich did not explicitly state that the switching fabric was an integrated switching fabric.

On the other hand, Prael did teach a cluster management engine communicably coupled to the plurality of nodes (column 1, lines 60-64) and operable to dynamically allocate a particular subset of the plurality of nodes to a particular job selected from a

queue comprising a plurality of jobs (column 1, lines 60-64 and column 7, lines 36-41) based on the determined status of each of one or more of the plurality of nodes and execute the job using the particular subset (column 7, lines 41-50, where the availability/unavailability of the nodes is the utilized status information). Prael also taught that the switching fabric was an integrated switching fabric (Column 4, lines 10-17). Both the systems of Blumrich and Prael are directed towards management for multiprocessing systems and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of Blumrich, to use integrated switching fabrics and dynamic job allocation, as taught by Prael. Integrated switching fabrics improve the efficiency of the system by providing more space within a system and also faster (integrated) switching. Furthermore, dynamic job allocation allows a system to make the most efficient use of its resources. Lastly, although not inherent, dynamic job allocation is given within a multi-processing system.

7. With respect to claim 12, Blumrich disclosed a method comprising: determining a status of each of at least a subset of a plurality of nodes ([0210], lines 4-9) , each node comprising a switching fabric integrated to a card and at least two processors integrated to the card (0057, lines 1-14); dynamically executing the job using the particular subset ([0226], lines 20-22, using run-time library information to configure hardware is dynamic allocation), but Blumrich did not explicitly state the fabric being a switching fabric or the dynamic allocation being particular subset of the plurality of nodes to a particular job selected from a queue comprising a plurality of jobs, based on the determined status of

each of one ore more of the plurality of nodes. However, Prael did teach the fabric being a switching fabric (Column 4, lines 10-17) and the dynamic allocation being particular subset of the plurality of nodes to a particular job selected from a queue comprising a plurality of jobs (column 1, lines 60-64 and column 7, lines 36-41) based on the determined status of each of one ore more of the plurality of nodes (column 7, lines 41-50).

Both the systems of Blumrich and Prael are directed towards management for multiprocessing systems and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of Blumrich, to use integrated switching fabrics and dynamic job allocation, as taught by Prael. Integrated switching fabrics improve the efficiency of the system by providing more space within a system and also faster (integrated) switching. Furthermore, dynamic job allocation allows a system to make the most efficient use of its resources. Lastly, although not inherent, dynamic job allocation is given within a multi-processing system.

8. As for claims 2, 13, and 24, the combination of Blumrich and Prael disclosed all of the limitations described in claims 1, 12, and 23 (respectively). In addition, Blumrich taught being further operable to determine a topology of the plurality of nodes ([0255], lines 1-5) based, at least in part, on the determined status of the nodes ([0225], lines 1-10, the network determines the status of elements to decide if it favors another method).

9. As for claims 3, 14, and 25, the combination of Blumrich and Prael disclosed all of the limitations described in claims 2, 13, and 24 (respectively). In addition, Blumrich taught wherein the topology comprises a three dimensional Torus ([0015], lines 1-5).

10. As for claims 34, 35, and 36, the combination of Blumrich and Prael disclosed all of the limitations described in claims 1, 12, and 23 (respectively). In addition, Blumrich taught wherein the card is a motherboard (0006, lines 1-10).

11. As for claim 9, the combination of Blumrich and Prael disclosed all of the limitations described in claim 1. In addition, Blumrich Prael taught wherein the queue comprises a plurality of jobs awaiting execution, each job submitted by a respective user; and the queue is on of a plurality of queues, each queue associated with a respective virtual cluster of nodes (column 1, lines 60-64 and column 7, lines 36-41).

12. Claims 4-5, 8, 10-11, 15-16,19-22, 26-27, and 30-33 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blumrich, in view of Prael, as applied to claims 1, 12, and 23 (respectively), and in view of Allen et al. (The Cactus Worm: Experiments with Dynamic Resource Discovery and Allocation in a Grid Environment), hereinafter Allen.

13. As for claims 4, 15, and 26, the combination of Blumrich and Prael taught all of the limitations described in claims 1, 12, and 23 (respectively). In addition, Blumrich taught wherein each node comprises at least one host channel adapter (figure 2, item

49, where an Infiniband link is comprised of host channel adapters); and the cluster management engine further operable to dynamically allocate a virtual cluster in the plurality of nodes. But Blumrich did not explicitly state a virtual cluster nor did he teach a dynamically allocated subset for executing the job comprising at least a subset of the virtual cluster. However, Allen taught such a system (Introduction, lines 4-8). It would have been obvious to a person of ordinary skill in the art at the time of the invention to modify the teachings of Blumrich in order to utilize dynamic allocation of virtual clusters, as taught by Allen, in order to make an efficient system that can adjust to different size workloads without undue hardship.

14. As for claims 5, 16, and 27, the combination of Blumrich and Prael disclosed all of the limitations described in claims 4, 15, and 26 (respectively). In addition, Blumrich taught the cluster management engine further operable to: dynamically allocate a second particular subset of nodes in the virtual cluster; and execute a second job using the second particular subset ([0024], lines 1-7; [0028], lines 3-7, where parallel processing is the second subset of nodes).

15. As for claims 8, 19, and 30, the combination of Blumrich and Prael disclosed all of the limitations described in claims 4, 15, and 26 (respectively). In addition, Blumrich taught the cluster management engine further operable to dynamically allocate a second cluster in the plurality of nodes ([0024], lines 1-7; [0028], lines 3-7, where

parallel processing is the second subset of nodes). But Blumrich did not explicitly state allocating virtual clusters. However, Allen did (Introduction, lines 4-8).

Both the systems of Blumrich and Allen are directed towards management for multiprocessing systems and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of Blumrich to allocate virtual clusters, as taught by Allen, in order to provide a system with increased efficiency, as the system will be able to perform small tasks along side of big tasks, instead of making all tasks wait for the larger task to be completed.

16. As for claims 20, and 31, the combination of Blumrich and Prael disclosed all of the limitations described in claims 19, and 30 (respectively). In addition, Blumrich taught the second virtual cluster comprises different nodes from the first virtual cluster ([0028], lines 3-11).

17. As for claims 10, 21, and 32, the combination of Blumrich and Prael disclosed all of the limitations described in claims 1, 12, and 23 (respectively). In addition, Blumrich taught wherein to execute the job using the particular subset (see the rejection for claim 1) the cluster management engine operable to: receive a job request comprising one or more parameters (in performing calculations, the process of receiving and carrying out a job is taking place); dynamically allocate the subset of the plurality of nodes, and execute the job using the dynamically allocated subset ([0226], lines 2-22). But Blumrich did not explicitly state determine dimensions of the job based, at least in part,

on the one or more job parameters or based at least in part, on the determined dimensions.

However, Allen did teach such a system (Introduction, lines 4-8, where changing based off of characteristics requires that dimensions are determined). It would have been obvious to a person of ordinary skill, in the art, at the time of the invention to modify the teachings of Blumrich in order to determine dimensions, as taught by Allen. Doing so provides a system that can better handle its workloads. Furthermore, determining dimensions or status would likely be used in dynamic allocation of any sort.

18. As for claim 11, 22, and 33, the combination of Blumrich and Prael disclosed all of the limitations described in claims 10, 21, and 32 (respectively). In addition, Blumrich taught the cluster management engine further operable to: select a policy ([0255], lines 1-15, where the Torus, global tree, etc are different policies) based on the job request; and dynamically determine the dimensions of the job further based on the selected policy ([0226], lines 20-22).

19. Claims 6-7, 17-18, and 28-29 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blumrich, in view of Prael, in view of Allen, as applied to claims 4, 15, and 26, and in further view of Zircher et al. (Pre-Grant Publication No. US 2003/0217105 A1), hereinafter Zircher.

20. As for claims 6, 17, and 28, they are rejected to on the same basis as claims 4, 15, and 26 (respectively) above. However, the combination of Blumrich, Prael, and Allen did not explicitly state the virtual cluster associated with a user group. However, Zircher did teach such a system ([0097], lines 10-13, where the designated devices form the user groups). It would have been obvious to a person having ordinary skill in the art at the time of the invention to modify the teachings of Blumrich, Prael, and Allen in order to utilize user groups, as taught by Zircher. Doing so greatly increases the autonomy of the s use and allowing it to serve a higher volume of customer/clients.

21. As for claims 7, 18, and 29, they are rejected on the same basis as claims 6, 17, and 28 (respectively) above. In addition, Zircher taught the cluster management engine further operable to verify a user submitting the job based, at least in part, on the user group ([0097], lines 10-13 and [0102], lines 1-2, in the process of utilizing the access control list, the verification of the user is taking place).

22. Claims 37-39 are rejected under 35 U.S.C. 103(a) as being unpatentable over Blumrich, in view of Prael, and in further view of Official Notice.

23. With respect to claims 37, 38, and 39, Blumrich disclosed Software residing at one or more computer systems collectively operable to execute the software (0024, lines 28-33 and 0025, lines 4-7), the software comprising: a plurality of cluster agents, each cluster agent associated with one of a plurality of nodes, the cluster agent

operable to determine a status of the associated node ([0057], lines 1-14), each node comprising: at least two first processors integrated to a first card and operable to communicate with each other via a direct link between them (0057, lines 3-4); the first processors communicably coupled to the first switch, the first switch operable to communicably couple the first processors to six or more second cards each comprising at least two second processors integrated to the second card and a second switch integrated to the second card operable to communicably couple the second processors to the first card and at least five third cards each comprising at least two third processors integrated to the third card and a third switch integrated to the third card (0203, lines 8-10 and figures 1 and 5);

Blumrich also disclosed the first processors being operable to communicate with particular second processors on a particular second card via the first switch and the second switch on the particular second card; the first processors being operable to communicate with particular third processors on a particular third card via the first switch, a particular second switch on a particular second card between the first card and the particular third card, and the third switch on the particular third card (0057, lines 1-14, where the first switch is the same as the second switch).

However, Blumrich did not explicitly state a cluster management engine communicably coupled to the plurality of nodes and operable to dynamically allocate a particular subset of the plurality of nodes to a particular job based on the determined status of each of one or more of the plurality of nodes and execute the job selected from

a queue comprising a plurality of jobs using the particular subset. Also, Blumrich did not explicitly state that the switching fabric was an integrated switching fabric.

On the other hand, Prael did teach a cluster management engine communicably coupled to the plurality of nodes (column 1, lines 60-64) and operable to dynamically allocate a particular subset of the plurality of nodes to a particular job selected from a queue comprising a plurality of jobs (column 1, lines 60-64 and column 7, lines 36-41) based on the determined status of each of one or more of the plurality of nodes and execute the job using the particular subset (column 7, lines 41-50, where the availability/unavailability of the nodes is the utilized status information). Prael also taught that the switching fabric was an integrated switching fabric (Column 4, lines 10-17). Both the systems of Blumrich and Prael are directed towards management for multiprocessing systems and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of Blumrich, to use integrated switching fabrics and dynamic job allocation, as taught by Prael. Integrated switching fabrics improve the efficiency of the system by providing more space within a system and also faster (integrated) switching. Furthermore, dynamic job allocation allows a system to make the most efficient use of its resources. Lastly, although not inherent, dynamic job allocation is given within a multi-processing system.

However, the teachings of the combination of Blumrich and Prael did not explicitly state two processors on a card that communicate with a central authority without communicating via either second processor on the particular second card. However, the examiner gives official notice that the purpose of multi-processing is to

improve the efficiency of processing tasks and therefore, not having the two processors communicate would require a third authority to designate the tasks, just increasing overhead of the system and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of the combination of Blumrich and Prael to utilize processors that don't co-communicate.

Response to Arguments

24. Applicant's arguments filed 04/28/2010 have been fully considered but they are not persuasive.

25. The applicant argues on page 12 that "***Blumrich*** **discloses that a separate, external 100 Mbps Ethernet switch connects a subset of computing nodes to each other and to an I/O node. (Fig. 5; ¶¶ 0057, 0066, 0079, 0146). In addition,** ***Blumrich*** **discloses that a separate, external Gigabit Ethernet or Infiniband switch connects the I/O node to an external RAID system. (Fig. 5; ¶¶ 0067, 0081). Thus,** ***Blumrich*** **discloses Ethernet and Infiniband switches that are separate from and external to the computing nodes and I/O nodes. Merely disclosing switches that are separate from and external to computing nodes and I/O nodes does not teach, suggest, or disclose "a plurality of nodes, each node comprising a switching fabric integrated to a card".** However, even assuming arguendo, claim 1 is rejected by Blumrich in view of Prael and Prael teaches multiprocessors. Likewise, multiprocessors contain a switching fabric between them. In order to provide a

reference for evidence, as requested by the applicant, view the abstract of Fu et al. (Patent No. US 6,633,945 B1).

26. The applicant also argues on pages 12 and 13 that "**the cited portion of *Prael* discloses "a collection of interconnected stand alone computers...working together as a single integrated computing resource." (Col. 4,11. 9-17). Merely disclosing stand alone computers (which *Prael* calls "nodes") that "work together" as a single integrated computing resource does not teach, suggest, or disclose that "each node compris[es] a switching fabric integrated to a card"**" and "**the Examiner appears to rely on *Prael,* stating "Prael (Column 4, lines 7-20) does directly state that a node can be 'a symmetric multiprocessor.' Thus, the nodes do posses a switching fabric as it is necessary for the communication between multiple processors on a single board." *Final Office Action* at 12-13. However, the Examiner has provided no evidence that this is true even now, let alone at the time of Applicants' invention**." Thus, in order to provide the requested evidence, Fu et al. (Patent No. US 6,633,945 B1) has been provided (see the abstract).

27. The applicant also argues on pages 14 and 15 that "**Merely mapping a computational task to available processors and, as characteristics change, extending to other resources, as disclosed in *Allen,* does not teach, suggest, or disclose "determin[ing] dimensions of the job" or "dynamically allocat[ing] the particular subset based, at least in part, on the determined dimensions"**" and

"**even assuming for the sake of argument only that the "characteristics of the computation" can be equated to the claimed "dimensions of the job" (which Applicants do not concede), *Allen* does not disclose any determination of such "characteristics of the computation".**" However, if Allen's system is able to make changes, based on the changes in the characteristics (abstract, lines 4-8), then it is inherent that the system be able to determine the changed characteristics.

28. The applicant also argues on page 16 that "the Examiner states, "**the purpose of multi-processing is to improve the efficiency of processing tasks and therefore, not having the two processors communicate would require a third authority to designate the tasks, just increasing overhead of the system and therefore, it would have been obvious to a person of ordinary skill in the art, at the time of the invention, to modify the teachings of the combination of *Blumrich and Fung* to utilize processors that don't co-communicate." *See Final Office Action* at 12. Applicants do not necessarily agree. For example, "increasing overhead of the system" may provide a reason that a person of ordinary skill in the art would not include a central authority.**" However, the argument that the examiner intended to make was that the technology was well known, but that it would have likely made the system less efficient; thus, they decided not to implement the technology.

Conclusion

29. Any inquiry concerning this communication or earlier communications from the examiner should be directed to JOSEPH GREENE whose telephone number is (571)270-3730. The examiner can normally be reached on Mon - Thu, 8:00AM - 4:00Pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Follansbee can be reached on 5712723964. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

JLG

/John Follansbee/

Supervisory Patent Examiner, Art Unit 2451